

Scalable replica-exchange framework for Wang–Landau sampling

Thomas Vogel,^{1,*} Ying Wai Li,² Thomas Wüst,³ and David P. Landau¹

¹*Center for Simulational Physics, The University of Georgia, Athens, GA 30602, USA*

²*National Center for Computational Sciences, Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA*

³*Scientific IT Services, ETH Zürich, 8092 Zürich, Switzerland*

We investigate a generic, parallel replica-exchange framework for Monte Carlo simulations based on the Wang–Landau method. To demonstrate its advantages and general applicability for massively parallel simulations of complex systems, we apply it to lattice spin models, the self-assembly process in amphiphilic solutions, and the adsorption of molecules on surfaces. While of general, current interest, the latter phenomena are challenging to study computationally because of multiple structural transitions occurring over a broad temperature range. We show how the parallel framework facilitates simulations of such processes and, without any loss of accuracy or precision, gives a significant speedup and allows for the study of much larger systems and much wider temperature ranges than possible with single-walker methods.

I. INTRODUCTION

In Monte Carlo simulations, one is interested in stochastically sampling the configurational space of a model system by the creation of a chain of consecutive microstates X_i (via a random ‘walker’):

$$X_0 \xrightarrow{\text{update}} X_1 \xrightarrow{\text{update}} \dots \xrightarrow{\text{update}} X_n.$$

The resultant ‘dynamics’ is artificial and depends upon the mechanism of proposing new microstates; n is usually a large number $\gg 10^6$. The proposed new microstate is accepted with a certain probability which determines the statistical ensemble. In statistical physics, most common simulations are carried out in the canonical ensemble by fixing the volume and particle density of a system and by setting the acceptance probabilities equal to the fraction of Boltzmann factors between the actual and the proposed new microstate [1].

A serious weakness of this scheme is well known: Typically, there are barriers in the free energy landscapes of complex systems, and the time it takes to overcome these barriers grows exponentially with their heights. Various Monte Carlo methods have been developed to confront the challenge of sampling such rough free energy landscapes [2–6] and to carry out simulations in an ensemble where the walker is not hindered by any barrier. For instance, Wang–Landau sampling [7, 8] has been shown to be very effective in overcoming energy barriers by iteratively determining the energy density of states (DOS) of a system and seeking to perform a random walk in energy space (‘flat histogram’); i.e. eventually performing a walk through configurational space such that all possible energies are visited uniformly. Wang–Landau (WL) sampling was successfully used in many scientific problems [9–22]. Another powerful approach is parallel

tempering or replica-exchange Monte Carlo [23–27]. Here the idea is to run multiple simulations in canonical ensembles at different temperatures and propose replica or conformational exchanges between them, which are accepted with a probability according to their respective Boltzmann weights. This configurational mixing among different walkers greatly alleviates the trapping problem near conformational or energy barriers.

Whereas a parallel implementation of replica-exchange Monte Carlo is straightforward, efficient and correct parallelization of the Wang–Landau (WL) algorithm has posed some subtle difficulties. Previous attempts have, for example, focused on running multiple, independent WL samplers which simultaneously update the same density of states [29–31]. However, a recent, massively parallel implementation of this approach [31] has revealed that inter-dependencies among the various WL walkers can introduce an erroneous bias in the estimate of the DOS and thus render this parallelization scheme highly problematic. Another WL parallelization, which is based on a simple splitting of the global energy range into smaller, independently sampled, non-overlapping energy windows, is also unsatisfactory. If the energy windows chosen are too small, configurational space may not be sampled correctly anymore due to ergodicity breaking, resulting again in subtle systematic errors (this effect becomes particularly pronounced in 2D WL simulations, where one samples a joint DOS, e.g., in energy and magnetization space [19]). Moreover, the total simulation time is bound to the WL convergence time of the slowest walker (generally in the low energy region).

Unlike replica-exchange Monte Carlo [24], such issues have severely limited the use of WL sampling as a means in large-scale parallel Monte Carlo simulations. A natural route towards successful parallelization of WL sampling would be to combine it with the benefits of replica-exchange Monte Carlo. Variants of this general idea have been used for specific applications by other authors on a small scale [20, 28, 35], however, no details about the implementation, parallelization, or the effect on the performance are provided. The critical issue of combining results from different replicas is omitted entirely. Our ap-

* thomasvogel@physast.uga.edu;

Present address: Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA

proach [32] is different by introducing a fast but generally applicable framework suitable for massive parallelization combined with a general and precise scheme to combine results from all walkers. We also start by running individual WL walkers in many overlapping energy windows [8, 33, 34] covering the whole energy range and allow for conformational exchanges between walkers according to the actual WL weights. In addition to splitting up the global energy range, we run multiple walkers within each energy window. In order to avoid any possible bias, these walkers are *independent* and fulfill convergence, or flatness criteria, individually. However, they work together by merging their weights after each Wang–Landau iteration, reducing the statistical error *during* the simulation.

In this paper we investigate this hierarchical parallel Wang–Landau scheme in detail with the aim to present a framework which nevertheless remains, in a sense, as easy to implement as the original single-walker WL method itself. Most importantly, the scheme does not depend on having previous knowledge about the system under investigation. We apply the method to a standard benchmark model in statistical mechanics, the 10-state Potts model, and to two cutting-edge problems which have attracted great interest recently: the self assembly of amphiphilic peptides into micelles and lipid bilayers, and the surface adsorption process of polymers and proteins. In Sec. II we introduce the complete parallel Wang–Landau framework in detail and in its most general formulation. We discuss the key points to make the framework efficient. In Sec. III we describe the models which will be used in Sec. IV to assess the applicability and accuracy of our parallel WL scheme (IV A). Furthermore, we present results of performance and scaling analyses in Secs. IV B and IV C. Finally, we combine our findings and suggest further potential, methodological improvements in Sec. V.

II. GENERAL PARALLEL FRAMEWORK FOR WANG–LANDAU SIMULATIONS

In generalized ensemble Monte Carlo simulations, one is interested in the density of states (DOS) $g(E)$ over a large energy range (E). In our framework, $g(E)$ will be determined in parallel by employing multiple computing cores, resulting in multiple, individual DOS pieces. If there is no generic and precise way of putting these pieces together in the post processing step of the simulation, the whole framework becomes meaningless. We will hence split this section into two parts: the production of the individual DOS pieces, and their assembly into a global density of states. Both are equally important aspects of our framework.

A. Replica-exchange Wang–Landau sampling

In the standard Wang–Landau method [7, 8, 33], a single walker (i.e., a chain of microstates) samples the con-




FIG. 1. Possible partitioning of the global energy range into nine windows: (a) the trivial case with an overlap of 100 %; (b) energy windows are slightly shifted with an overlap of 90 %; (c) an example of a run-time balanced splitting with an overlap of (at least) 75 % (note that the widths of the windows vary for different systems). Several walkers might run in each window.

formational space in an energy range between E_{\min} and E_{\max} , improving the estimate of $g(E)$ iteratively. The microstates are sampled according to the actual weights $1/g(E)$, which are adapted on the fly in the following way: After each Monte Carlo trial, the estimator for $g(E(X))$ for the walker residing at microstate X is multiplied by a modification factor f . At the same time the histogram of visited energies $H(E(X))$ is incremented. Whenever there are ‘sufficient’ entries in $H(E)$, the histogram will be reset and the modification factor will be decreased, for example as $\ln f \rightarrow \ln f/2.0$. Initially, $g(E) = 1.0$ and $H(E) = 0, \forall E$, and $\ln f = 1.0$. There are different approaches to ensure that there are ‘sufficient’ entries in the histogram for all values of E , or for all energy bins, respectively. In the original formulation [7], the histogram $H(E)$ is required to be ‘flat’, i.e. no value for $H(E)$ may be smaller than a certain percentage of the average histogram value. Other authors [36] only insist that there are more entries in the histogram for each E than a minimal threshold number which increases with the actual value of f . Additionally, a method has been proposed where the modification factor f decreases as $1/t$, where t is the simulation time [37]. (We emphasize that our parallel framework presented here is general in the sense that it does not depend on such details.) The estimated density of states converges to the true one with increasing number of iterations, and the simulation is terminated when the modification factor reaches a minimal

value f_{\min} , typically set to $\ln f_{\min} < 10^{-6} - 10^{-8}$. For all practical purposes, the WL walker eventually performs a random walk in energy space.

For large systems, this very efficient generalized-ensemble sampling can be enhanced by making use of multiple processors working in parallel. This can be done, for example, by splitting up the global energy range into smaller energy windows and estimating the density of states for the respective windows by independent walkers [8, 33, 34]. See Fig. 1 for three examples of possible energy windows. Following the general idea of replica-exchange methods [23, 24, 27], it is natural to allow for replica exchanges between independent walkers in WL sampling as well if the energy windows overlap. From the detailed-balance condition for the combined trial move, the acceptance probability P_{acc} is derived for the exchange of conformations X and Y between walkers i and j :

$$P_{\text{acc}} = \min \left[1, \frac{g_i(E(X)) g_j(E(Y))}{g_i(E(Y)) g_j(E(X))} \right], \quad (1)$$

where $g_i(E(X))$ is the current estimator for the density of states of walker i at the energy of its present conformation (microstate X).

For replica-exchange attempts, the walker synchronously communicates with a neighbor in an adjacent energy window, see Fig. 2 for an illustration. (Note, in principle one could allow replica exchanges between walkers in any windows with non-zero energy overlap.) In addition to replica exchange, we will also run multiple, independent walkers in each individual energy window. Each walker carries its own estimator for the density of states $g(E)$ and its own histogram $H(E)$ and is required to fulfill the ‘flatness’ criterion individually. In particular, individual histograms $H(E)$ will not add, a feature which eliminates the potential for systematic errors as observed previously [31]. However, walkers in the same energy window will merge and average their individual $g(E)$ estimator before simultaneously proceeding to the next Wang–Landau iteration step. This averaging results in a reduction of systematic errors of $g(E)$ during the course of the simulation and, thus, reduces the overall convergence time. The parallel simulation terminates when every walker has reached the final modification factor f_{\min} . Walkers that have converged to the final modification factor early will continue walking in order to allow for replica exchange with walkers which have not finished yet.

As in any replica-exchange scheme, one has to ensure that the exchange of replica, and thus the overall flow of the simulation, does not get stuck; and, hence, each replica can perform frequent round-trips over the entire energy range. Obviously, there must be some overlap between communicating energy windows in the first place. Furthermore, in order to make this framework efficient, two questions arise which have to be answered individually for different systems:




FIG. 2. The communication for the replica exchange between energy windows alternates between the one marked by red (left) and the one marked by blue (right) arrows.

(i) What is a reasonable overlap between neighboring energy windows? Certainly, the overlap must not be too small or it would lead to small acceptance rates for conformational exchanges, just like in parallel tempering simulations when two probability distributions have a very small overlap. On the other hand, excessive overlaps will make the scheme inefficient, as in the extreme case one would have n walkers in energy windows identical to the global energy range, see Fig. 1 a.

(ii) What is a good number of energy windows? Generally, one would like as many windows as possible to improve the scaling of the simulation and exploit the capabilities of modern parallel computers. On the other hand, when the number of energy windows increases, the time for complete round trips also increases. Also, given a fixed global energy range and equal-size energy windows, the number of energy windows directly relates to their width, and too small energy windows will eventually increase the occurrence of systematic errors, even though the risk of ‘locking out’ of parts of the conformational space practically becomes negligible due to the replica-exchange mechanism. In any case, this problem obviously depends much more on the actual system than the first question does.

On the algorithmic level, the scaling of the framework is important for its general applicability and has to be investigated:

(iii) How does the performance depend on the number of computing cores used? In the optimal case, a method would achieve both strong scaling and weak scaling, i.e., if more cores are applied, one wants to be able to obtain results faster or to simulate larger systems, respectively, or even both. In that context it is important to know how the number of Wang–Landau walkers within each energy window affects the performance of the framework and the error of the final DOS estimate.

With the help of different ‘test’ systems, ranging from discrete lattice spin models to coarse-grained molecular systems in continuous space, we will examine these questions in more detail and demonstrate the potential of our novel approach. All models are qualitatively different and known to show complex behavior. We will use the following notations throughout this paper:

h : number of energy windows,

m : number of independent walkers per energy window,

o : overlap with the next lowest energy window ($0 \leq o \leq 100\%$),

$s_h(o)$: speedup of the parallel simulations with overlap o and constant h (see Eq. (7) for definition),

$s_o(h)$: speedup of the parallel simulations with h energy windows at constant overlap o .

B. Concatenation of DOS pieces

At the end of a parallel WL simulation we are left with h DOS pieces which first need to be put together carefully before any thermodynamic quantities, such as the internal energy or specific heat, can be calculated. This DOS concatenation procedure is a delicate technical challenge since even small artificial steps or kinks in the entropy ($\propto \ln[g(E)]$) may cause significant artificial peaks and oscillations in observables like the specific heat. Therefore, simply joining DOS pieces at some fixed positions likely results in discontinuities in the DOS, which, in turn, cause significant artifacts in derived quantities. To minimize such erroneous effects, we proceed as follows:

(i) We calculate the first derivative $\partial \ln(g(E))/\partial E$ (i.e. the inverse microcanonical temperature) for each DOS piece in the overlap region between two adjacent energy windows. A large number of overlap data points allows us to use higher order 5-point approximations of the derivative with varying step width of the order of up to 10 (depending on the system and the width of the energy windows) which provide very smooth estimates for the derivatives.

(ii) We determine the point where the inverse microcanonical temperatures of the two overlapping DOS pieces coincide the best, then connect the DOS pieces at this point and cut the ‘overhanging tails’ at the respective sides of each piece, thus avoiding non-differentiable points in the resulting entropy by construction. This technique of connecting DOS pieces turned out to be indispensable for obtaining accurate results as shown below; for more details, see [38, 39].

For a rigorous error analysis, we perform n independent parallel simulations from which we get $n \times h$ individual DOS pieces. We then calculate the mean of the n DOS pieces for each energy window before proceeding to steps (i) and (ii) above. This yields the mean global DOS. To estimate statistical errors, we apply a bootstrap resampling technique [40]. That is, we randomly choose n DOS pieces (with repetitions) for each energy window and apply the above two steps. This procedure is repeated multiple times (e.g. 200) yielding multiple re-sampled global DOS. From these we calculate the statistical errors of the global DOS and its derived observables. The entire technique is unambiguous and very precise; moreover, it has an advantage that due to the random selection of DOS pieces during the bootstrap analysis,

connection points are always at different positions leading to statistical smoothing of potentially remaining artifacts in derived quantities.

III. MODELS AND MODEL-SPECIFIC ALGORITHMIC DETAILS

A. The Potts model for lattice spins

The first model we use to test our parallel WL scheme is the well studied Q -state Potts model for lattice spins in two dimensions [41, 42]. It is a common test bed for novel simulation methods (see [3, 8, 43] for examples), as the system size is scalable in a straightforward manner and exact results exist, e.g. for the infinite size transition temperature of the first-order phase transition between the ordered and disordered phases for $Q > 4$ [44]. The Hamiltonian is given by

$$\mathcal{H} = - \sum_{\langle i,j \rangle} \delta(q_i, q_j), \quad (2)$$

where the spins q_i can take values $q_i = 1 \dots Q$ and the sum is over all nearest neighbor pairs $\langle i, j \rangle$. For this study, we choose the 10-state Potts model, i.e., $Q = 10$, and use periodic boundaries. The total energy range is given by $-2N \leq E \leq 0$, where $N = L \times L$ is the total number of spins and L is the lattice size. As we use this model for demonstration purposes only, we perform the simplest Monte Carlo update move, a random, single spin update trial.

B. The hydrophobic-polar (HP) model for protein adsorption

The HP model [45] is a minimalist, coarse-grained lattice model used to study generic protein folding behavior. It classifies amino acids into only two types of monomers according to their affinity to water: hydrophobic (H) and polar (P). There is only an interaction between non-bonded, hydrophobic monomers occupying nearest-neighbor sites, with a coupling strength ε_{HH} .

Protein adsorption can be simulated with a slight modification to the HP model [46]. On a three-dimensional cubic lattice, a substrate, placed at the $z = 0$ plane, attracts H- and P-monomers in the HP chain with a strength ε_{SH} and ε_{SP} , respectively. The energy function of this model can then be written as:

$$E = -\varepsilon_{HH} n_{HH} - \varepsilon_{SH} n_{SH} - \varepsilon_{SP} n_{SP}, \quad (3)$$

where n_{HH} is the number of H-H interacting pairs and $n_{S[HP]}$ is the number of H- or P-monomers adjacent to the substrate. While periodic boundary conditions are imposed for the x - and y -directions, a non-attractive wall is placed at $z = N + 1$ to confine the simulation box from above, where N is the chain length of the sequence.




FIG. 3. Illustration of a partly adsorbed protein composed of 36 monomers in the HP model. Besides the surface contacts (not shown), there are two contacts between non-bonded hydrophobic monomers (indicated by arrows).

In the following, we consider a 36mer ($P_3H_2P_2H_2P_5H_7P_2H_2P_4H_2P_2HP_2$) [47] interacting with a weakly attractive surface by setting $\varepsilon_{HH} = 12$, $\varepsilon_{S[HP]} = 1$. This system has a very rugged density of states (see below) which makes it an excellent test case for our new parallel framework. Conformational updates are proposed by means of pull moves and bond-rebridging [22, 48–51]. For a fair comparison with previous results [52–54], pull moves make up 20% of the Monte Carlo moves, while bond-rebridging moves make up 80% of them. For the WL sampling, we use the originally introduced 80%-flatness criterion [7] and the simulations end when the modification factor reaches $\ln(f_{\min}) = 10^{-8}$. Fig. 3 shows a schematic part of the system, for more details see [52–54].

C. Generic model for amphiphilic molecules in solution

Lastly, we use a continuous, molecular model which is based on generic models previously employed to study the self-assembly of molecules into simple membranes, bilayers, or micelles. It includes small amphiphilic molecules surrounded by explicit solvent particles [55, 56]. There are three different types of coarse-grained particles: polar (P) and hydrophobic (H) monomers and the water or solution (W) particles. The amphiphilic molecules are composed of a polar head and two hydrophobic tail monomers: (P–H–H). The interaction between solution particles as well as between head monomers and tail monomers is purely repulsive and is modeled by a repulsive soft-core potential

$$U_{\text{repulsive}}^{\text{soft core}} = 4 \varepsilon_{\text{rep}} \left(\frac{\sigma_{\text{rep}}}{r_{ij}} \right)^9, \quad (4)$$

where $\varepsilon_{\text{rep}} = 1.0$ and $\sigma_{\text{rep}} = 1.05 \sigma$ (see below), following [55, 56]. All other non-bonded interactions are of Lennard-Jones type

$$U_{\text{LJ}} = 4 \varepsilon_{X-Y} \left[\left(\frac{\sigma_{X-Y}}{r_{ij}} \right)^{12} - \left(\frac{\sigma_{X-Y}}{r_{ij}} \right)^6 \right]. \quad (5)$$




FIG. 4. Illustration of particles and interaction parameters in the generic coarse-grained lipid model. P (blue): polar head monomer, H (yellow): hydrophobic tail monomer, W (red): solution particle ('water molecule').

r_{ij} is the Euclidean distance between two non-bonded particles i and j , the notation $X-Y$ stands for the interaction between particles of type W–W, W–P, P–P, and H–H, cf. Fig. 4. In principle, the parameters ε_{X-Y} and σ_{X-Y} can have different values for different $X-Y$ combinations allowing for the introduction of different energy scales. As we focus here on the technical aspects of our work, we fix all $\varepsilon_{X-Y} = 1.0$ and $\sigma_{X-Y} = \sigma = 2^{-1/6} r_0$, where $r_0 = 1.0$ defines the length scale in the system. The potentials in Eqs. (4) and (5) are cut off at $r_c = 2.5\sigma$ and shifted such that there are no discontinuities at this point.

Bonds are modeled using the finitely extensible, non-linear elastic (FENE) potential

$$U_{\text{bond}}^{\text{FENE}} = -KR^2 \ln \left[1 - \left(\frac{r_{i,i+1} - r_0}{R} \right)^2 \right]^{1/2}, \quad (6)$$

where $r_{i,i+1}$ is the length of a particular bond, $R = 0.3$ is half the width of the potential (clearly, the potential diverges at $r_{i,i+1} = r_0 \pm R$) and the 'spring' constant is set to $K = 40$. The equilibrium length r_0 at which $U_{\text{bond}}^{\text{FENE}} = 0$ coincides with the equilibrium length for the non-bonded potential U_{LJ} . Periodic boundary conditions apply.

Figure 5 shows examples of typical configurations for this model. The sequence of pictures shows snapshots of a system containing $M = 125$ amphiphilic molecules and $N = 1000$ particles in total at a number density of $\rho = 0.8$. One Monte Carlo (MC) sweep consists of N individual MC steps. Among these N steps are, on average, $3 \times M/10$ reptation moves, which we found to be essential in order to thoroughly examine the conformation space for large systems. The other moves are local displacements of individual particles. For the reptation move we first select, at random, a solution particle in the vicinity of one end of an amphiphilic molecule. This particle is then converted into either the new head or the new tail of the molecule, and the opposite-end monomer is converted into a solution particle. The bias introduced due to different numbers of neighbors at opposite ends of an amphiphilic molecule is accounted for in the calculation of the acceptance probability of this move; for




FIG. 5. Conformations of a system containing $M = 125$ amphiphilic molecules and a total of $N = 1000$ particles. (a) Random configuration, $E \approx -200$ (typical initial configuration for simulation). (b) Amphiphilic molecules assemble and form loose clusters, $E \approx -4000$. (c) Low-energy, single-cluster configurations with compact shape, $E \approx -5100$. Upper row: Only the amphiphilic molecules are shown, different colors mark different clusters of amphiphilic molecules. Periodic copies of the simulated system are shown. Bottom row: Simulation boxes including all particles.

more details, see [38]. For this model, we use the originally proposed ‘80-percent’ WL flatness criterion, the final modification factor is $\ln f_{\min} = 10^{-7}$.

IV. RESULTS AND DISCUSSION: ACCURACY AND PERFORMANCE

A. Number and size of energy windows; degree of overlap

When applying our parallel framework to the 10-state Potts model, (which shows a strong, temperature-driven first-order transition), we first vary the system size ($N = 1000, 4000, 9000$, and $16\,000$ spins) while keeping the energy window size ($\Delta E = 1000$) and overlap ($o = 75\%$) fixed. Consequently, the number of energy windows needed to cover the whole energy range increases. We used up to nine walkers in each energy window so that, in total, $\gtrsim 10,000$ cores were used for the biggest system. In Fig. 6a we plot the densities of states for the four lattice sizes. The data are composed of 77, 317, 717, and 1277 pieces for the $N = 1000, 4000, 9000$, and $16\,000$ spin systems, respectively; the density of states of the largest system covers more than $150\,000$ orders of magnitude. To further demonstrate the simulation challenge we show in Fig. 6b that transition states are suppressed by a factor of 10^{-17} – 10^{-18} for the largest system. To verify the results, we reproduced and extended the corresponding analysis from the original Wang–Landau paper [8]. Figure 6c shows the specific heats for all system sizes and Fig. 6d the finite size scaling analysis of the transition temperature. By extrapolating the peak position of the

specific heat peaks, we estimate the transition temperature in the thermodynamic limit to be $0.7012313(5)$ in agreement with the exact value 0.7012316 [44]. While this analysis would take years for a serial, single-walker code (the simulation for the 100×100 system for a single WL walker still takes several days), we obtained all results within a few hours applying our parallel scheme. The scaling analysis is presented below.

After this first ‘proof of concept’, we now apply our method to much more complex molecular models. Before presenting the physical results, we demonstrate the influence of the overlap of neighboring energy windows on the speedup $s_h(o)$ (as compared to single-walker runs) and acceptance rates for the replica exchange using the continuous model for amphiphilic molecules as described in Sec. III C. The system setup is the same as the one used in Fig. 5. We split the whole energy range of interest into nine windows ($h = 9$) and, for simplicity, employ only one walker in each window ($m = 1$). The overlaps take the following values: $o = 50\%, 62.5\%, 75\%, 87.5\%$ or 100% . Consequently, for $h = 9$, the widths of the individual energy windows for these overlap values are $1/5, 1/4, 1/3, 1/2$ (and 1) of the width of the global energy range, respectively, cf. Fig. 2 for the case ($h = 9, o = 75\%$). The speedup $s_h(o)$ is measured by the total number of MC steps (‘MC time’) it takes for the ‘slowest’ WL walker to satisfy the termination criterion ($t_{\text{term}}^{\text{parallel}}(h, o)$), compared to that for a single-walker WL simulation ($t_{\text{term}}^{\text{single}}$):

$$s_h(o) = \frac{t_{\text{term}}^{\text{single}}}{t_{\text{term}}^{\text{parallel}}(h, o)}. \quad (7)$$

The acceptance rate α is the percentage of accepted

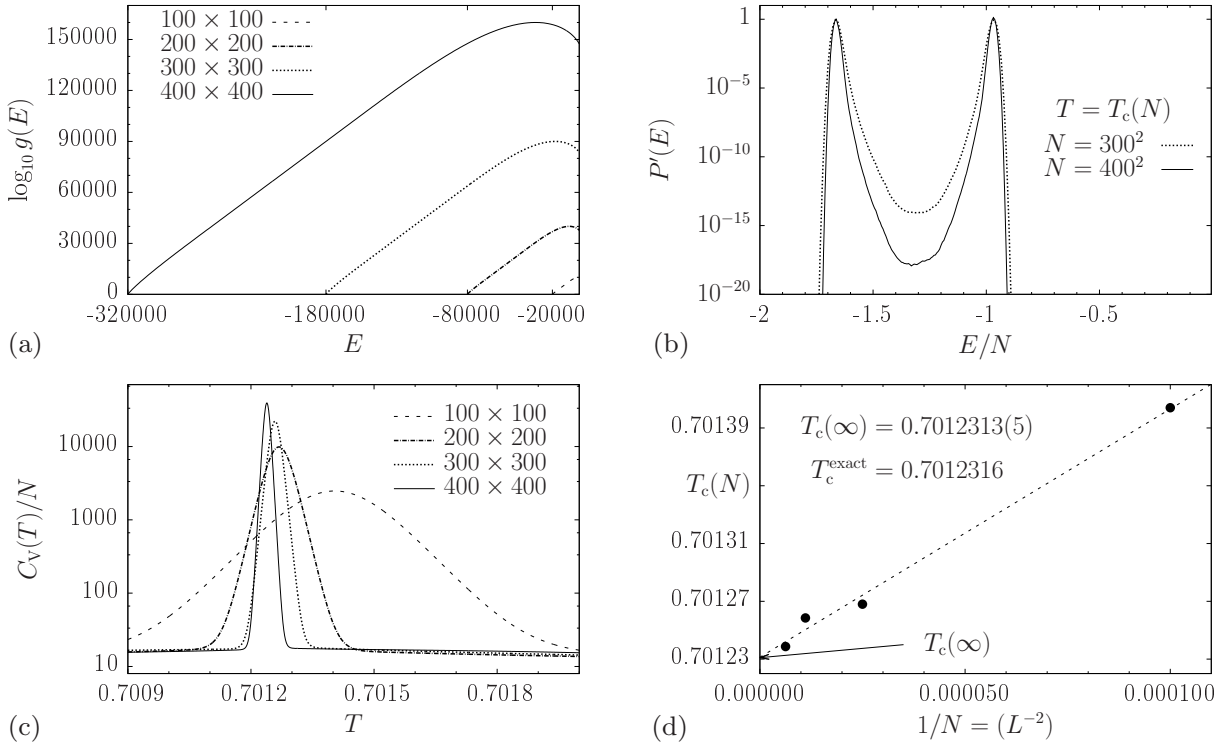


FIG. 6. a) Density of states of the 2D 10-state Potts model for four lattice sizes. The data are composed of 77 (100×100), 317 (200×200), 717 (300×300), and 1277 (400×400) pieces. b) Canonical distributions at the finite-size transition temperature for the larger systems. c) and d) The corresponding specific heat curves and the finite size scaling of the size-dependent transition temperature. Temperatures in (d) were obtained from peak positions in the specific heats in (c).

replica exchanges compared to the number of proposed exchanges.

The results of this test are shown in Table I. The statistical error of $s_h(o)$ is relatively large due to the well-known fact that the run time of traditional WL iterations at very small values of f can vary significantly for independent runs. For α , minimal and maximal values are given as the acceptance rates vary for different energy windows. Overall, the data allow us to draw the following consistent conclusions: An overlap of 100% does not improve performance. In fact, there is not much difference compared to n non-communicating single-walker simulations in this setup. However, the speedup is already significant for a relatively large overlap of 75% and increases rather slowly as the overlap decreases. On the other hand, even though acceptance rates for replica exchange in the range of 15–25% are satisfactory, the rates for $o = 75\%$ are much better and are comparable to the rates desired for canonical replica-exchange simulations. Keeping in mind that more overlap between neighboring energy windows will allow us to better connect the individual DOS parts later, we thus decided to fix the overlap to $o = 75\%$.

Related to these considerations is the problem of finding a reasonable number of energy windows. Given a fixed, global energy range as well as a fixed overlap, the number of energy windows defines their widths. As men-

o	50%	62.5%	75%	87.5%	100%
$s_h(o)$	5.5 ± 0.6	4.4 ± 0.5	4.2 ± 0.5	2.0 ± 0.3	1.0 ± 0.2
α	15–25%	25–40%	30–55%	35–75%	25–35%

TABLE I. Speedups ($s_h(o)$) and acceptance rates (α) for replica exchange at different overlaps (o) for the amphiphilic test system. $h = 9$, $m = 1$. See text for details.

tioned before, there are conflicting goals. On one hand, one would like to maximize the number of energy windows, however, individual windows must not be too small or systematic errors are induced by restricted sampling. In order to get reliable results, it is necessary that each replica performs walks through the *entire* energy range, which also means that each sample walks back and forth through *every* energy window. While the round trip times might be shorter when also allowing for replica exchange between next-nearest neighboring windows, there would be much more technical overhead and the implicit synchronization of the communication pattern would be much more complex.

For the amphiphilic test system, we find that h of the order of 10 leads to a good performance with respect to both speedup and round trip times for individual replica. Eventually, for the given system and a global energy range $E \in [-4500, -2000]$, we find that a split-




FIG. 7. Path of a replica through energy space (a) and energy windows (b). The samples perform complete round-trips within less than 5×10^6 MC sweeps. A conformational exchange between walkers is proposed every 10^4 sweeps, with acceptance rates between 30 and 55 % (cf. Table I). Grid lines in (a) correspond to the borders of the individual energy windows.

ting into nine energy windows with an overlap of 75% is a reasonable choice, cf. Fig. 2. In Fig. 7 we visualize the walk of a replica through the energy space (Fig. 7 a) and through the energy windows (b) during $\approx 2 \times 10^7$ MC sweeps (one MC sweep equals N MC steps, where N is the system size, i.e., the total number of particles). The replica performs a smooth walk through the energy space and completes a round trip approximately every 5×10^6 MC sweeps (a replica exchange between walkers was proposed every 10^4 MC sweeps). We confirmed that all replicas behaved similarly.

Figure 8 shows the logarithm of the DOS for both molecular systems: for the continuous-energy lipid model (cf. Fig. 8 a) the global energy range is chosen to be accessible by a single-walker WL simulation (run time of approximately one week). For the HP lattice protein (cf. Fig. 8 b) the complete energy range ($E \in [-241, 0]$) is sampled. Solid lines show data obtained from nine independent single-walker runs as a reference. The statistical error bars (estimated by the sample standard deviation σ) are smaller than the line thickness and are shown separately. The filled dots represent data obtained from a single parallel run with $o = 75\%$ and $h = 9$ for equal splitting energy windows. The absolute difference between the results from single-walker runs and the parallel run, Δ , is compared to the statistical errors from

the reference (single walker) runs. For both systems, we found that $\Delta \leq \sigma$ for practically all energies, i.e. the results from the parallel run are clearly within the error bars of the reference runs. The speedup measured for the lattice model was comparable to those shown in Table I: $s_9(75\%) = 4.3 \pm 0.4$.

Sec. IIB discussed the need for a technique to connect DOS pieces without introducing jumps or kinks (i.e., non-differentiable points) into the entropy or its derivatives. In particular, the precise results shown in Fig. 8 could not be obtained otherwise, and kinks in the density of states would inevitably result in artificial peaks in the heat capacity. To further illuminate the power of our framework, in Fig. 9 we show parts of the heat capacities corresponding to the density of states shown in Fig. 8. The data clearly show that the heat capacities obtained from the parallel and single walker (reference) runs are within mutual error bars. For the system of amphiphilic molecules (Fig. 9 a), the small, but particularly interesting, peak shown corresponds to the alignment of amphiphilic molecules during the transition from cylindrical structures to liquid bilayer sections. (The nature of the transition was unveiled during production runs measuring the distributions of a bond-orientation related order parameter [57] and the asphericity and prolateness resulting from the gyration tensor [58, 59]). We emphasize that this particular peak is hard to resolve in the heat capacity as it is almost overwhelmed by stronger signals at lower temperatures (not shown). The procedure to connect the individual DOS pieces is thus essential to resolve and separate the signal from artifacts and statistical noise. For further details and a study of the physics and thermodynamic behavior of the lipid self-assembly and transitions between different bilayer phases, see [60]. In Fig. 9 b we show the peak corresponding to the adsorption of the 36mer HP-protein to the substrate, which lies well below the collapse transition. See [52, 53] for more results and discussion.

Finally, Fig. 10 shows $\log g(E)$ for the amphiphilic system over a larger global energy range ($E \in [-5000, -2000]$), which is *not* accessible by single-walker simulations anymore. The data were obtained by a parallel sampling scheme with $h = 9$, $o = 75\%$ and $m = 3$, hence $n = 27$ walkers in total (a sketch of this setup is given below the data). As it is not possible to sample this energy range with a single walker, we can not measure the speedup for this case; however, we are able to cover an additional range of the density of states of several hundred orders of magnitude at low temperatures (marked by dotted lines in the plot). Owing to this extended sampling range it has been possible to uncover the intricacies of the low temperature lipid bilayer phases such as the liquid vs. gel phase – processes which are of physiological importance and are often studied in biochemistry and related fields; see [60] for further discussions. In the inset of Fig. 10, we give the deviation Δ of the results from the reference simulations in the energy range accessible for such.

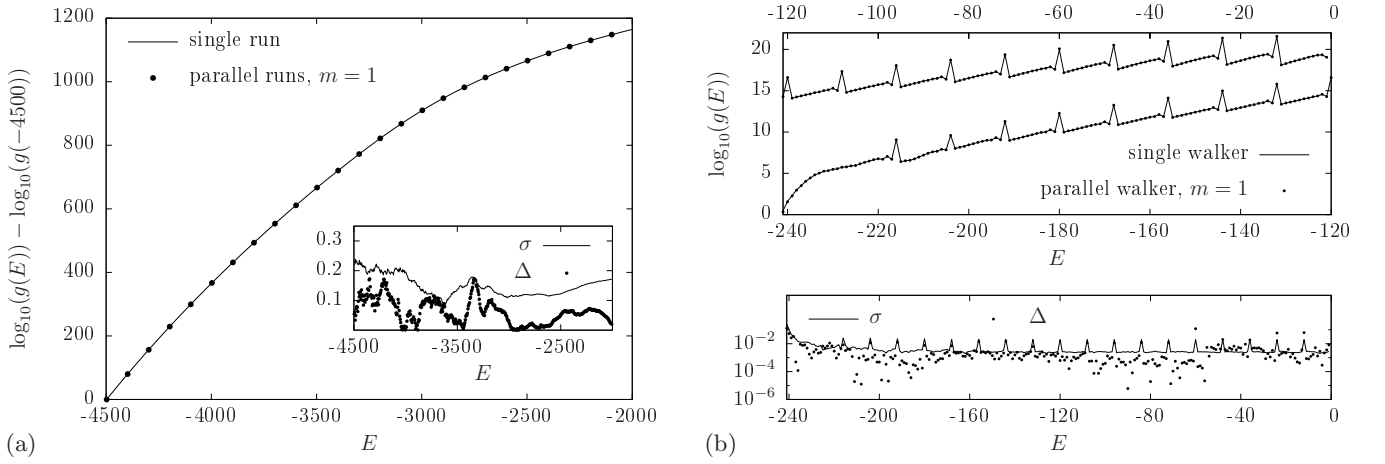


FIG. 8. Logarithm of the density of states obtained by single-walker Wang–Landau runs (solid line) and by a parallel run using nine energy windows with an overlap of 75% (dots). (a) Data for the system of amphiphilic molecules in solution. For clarity, only a small subset of data points is shown. (b) Data for the lattice HP 36mer on a weakly attractive surface. For clarity, the data are split into two curves corresponding to the lower and upper halves of the total energy range, respectively. The inset in (a) and the lower plot in (b) – note the logarithmic scale – illustrate the accuracy of the method. Solid lines show the standard deviation σ obtained from the serial, single-walker runs; dots show the absolute numerical difference Δ between data obtained by the single-walker runs and the parallel runs. All results obtained from the parallel runs are within the error bars of the reference runs.

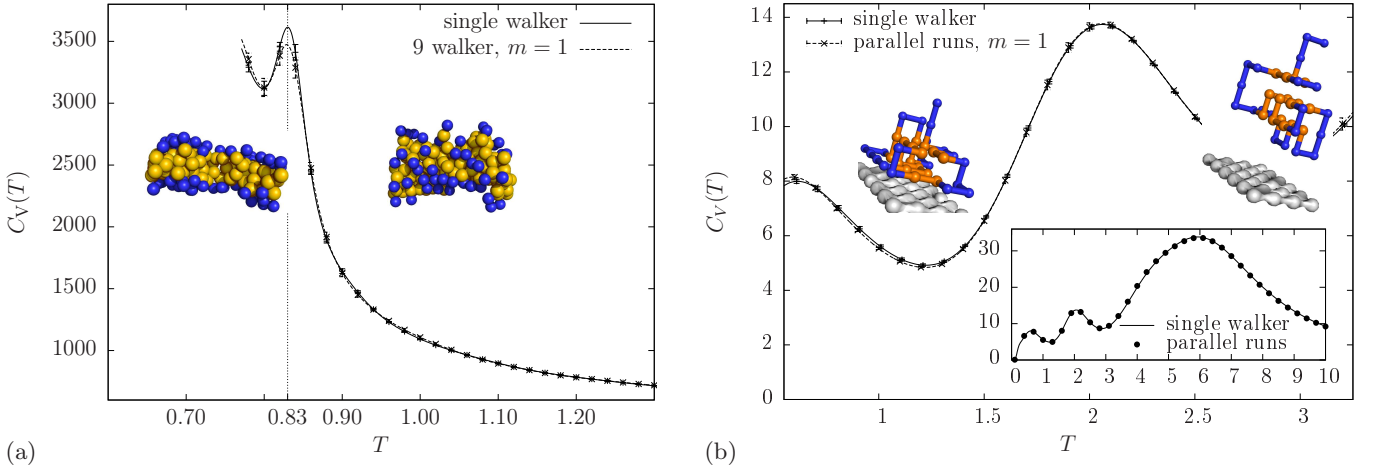


FIG. 9. Comparison of heat capacities obtained from single-walker runs and those from parallel runs. (a) The amphiphilic system in the region of the peak corresponding to the transition from cylindrical conformations into bilayer sections involving the alignment of amphiphilic molecules (for clarity, solution particles are not shown in pictures). (b) The 36mer in the HP model in the region of the adsorption peak. Pictures show representative adsorbed and desorbed conformations. The inset shows the same data in the full temperature range for reference. In both plots, the curves are within mutual error bars.

B. Multiple walkers per energy window

To quantify the effect of m walkers per energy window, we calculate the estimator of the error made after the k th WL iteration, which is denoted by $\Delta H_{m,k}$. We define this measure analogously to the case for single-walker WL simulations introduced in [61]. In Fig. 11 we plot $\Delta H_{1,k=25} / \Delta H_{m,k=25}(\sqrt{m})$, i.e., the error reduction after convergence of the WL scheme as a function of m . This is measured in the highest energy window of a run for the amphiphilic molecules system (cf. Fig. 10 for the

setup). The figure shows that the error reduces with \sqrt{m} , i.e., as for uncorrelated WL simulations, cf. [36]. This behavior is independent of k for iterations with $\ln(f) \lesssim 10^{-2}$, i.e. the error is reduced after each iteration during the simulation. Furthermore, increasing m can improve the convergence of the WL procedure by reducing the risk of statistical outliers in $g(E)$, which typically slows down subsequent iterations.

The use of multiple walkers within each energy window hence provides the possibility to complete WL iterations faster by choosing a weaker flatness criterion. For single-




FIG. 10. Logarithm of the density of states for a system containing $M = 75$ amphiphilic molecules and a total of $N = 1000$ particles. Comparison of results from single-walker runs with $E_{\min} = -4500$ (solid line, lower limits marked by dotted lines) and a parallel run with $E_{\min} = -5000$, $m = 3$ (dots). The inset shows the deviation Δ between both data. The diagram below the plot shows the parallel setup employed.

walker simulations this usually results in faster convergence but also larger statistical errors, but this can now be compensated for by applying multiple walkers to reduce the error. We can, for example, apply the flatness criterion proposed in [36]. Then, each walker leaves its modification factor unchanged until *all* walkers have accumulated a minimum number of histogram entries for each energy (bin), i.e., the flatness criterion is fulfilled if

$$H(E) \geq a/\sqrt{\ln f}, \quad \forall E \quad (8)$$

and for all m walkers inside the energy window. We thus guarantee, that each energy has been visited by independent walkers at least $(ma)/\sqrt{\ln f}$ times in total dur-




FIG. 11. Reduction of the WL error for different numbers of walkers per energy window, (m) after convergence ($k = 25$) of the WL iteration. The straight dotted line is only a guide to the eye (not a fit to the data).

ing each WL iteration, yet every walker still fulfills the flatness criterion independently. This ensures that systematic errors as found in [31] cannot occur. These m walkers will then merge their DOS estimator and proceed together to the next WL iteration. For every choice of a , it should be, in principle, possible to find a value for m such that the resulting statistical error is of the same order as for a single walker fulfilling a stricter flatness criterion. We applied this idea to simulations of the amphiphilic model setting the parameter a in Eq. (8) as small as 1 and measured a speedup of an order of 10. A *quantitative* number, however, can not be given as comprehensive simulations assuring that the statistical errors are of equal size for both approaches would require unreasonable computational costs just for this purpose. (This approach shares basic ideas with a recently proposed technique of merging histograms in multicanonical simulations [62].)

A setup with multiple walkers per energy window is also suitable for fault tolerant implementations [63] of the simulation code. Besides the fact that hardware or network failures otherwise usually result in a complete abortion of the program, a loss of some walkers due to such failures can easily be tolerated if there are multiple walkers in each window. In most cases, it would not ‘disconnect’ the ‘communication’ between the outer energy windows but only affect the statistical error in that region where walkers fail.

C. Strong and weak scaling behavior

We now consider the scaling of our method with increasing number of computing cores including a detailed analysis of both the speedup (strong scaling) and the potential for simulating larger system size (weak scaling) resulting from the increase of the number of individual energy windows. In Fig. 12a we show the speedup $s_{75\%}(h)$ for different numbers of energy windows $h \lesssim 15$ in the first WL iteration while keeping the number of processes in each energy window fixed. We employ $m = 1$ walker per energy window and use the continuous, amphiphilic system for this test. We find that the speedup scales linearly with the number of energy windows in this region, whereas the slope depends on the method of splitting of the energy range. With a run-time balanced speedup [32] we can achieve a slope > 1 , i.e., a speedup which is greater than the increase in the number of cores used.

Note again that for our purpose we are using a particular definition of speedup, which might differ from others commonly used, and that this surprising result can be attributed to a combination of two effects. First, due to the way the WL algorithm behaves when building up the entropy estimator (see, e.g., [36]), the total effort (e.g. MC sweeps) needed for multiple walkers to fill up their individual histograms can, depending on the energy splitting and the shape of the entropy curve, be less than the work needed for a single walker to fill a histogram

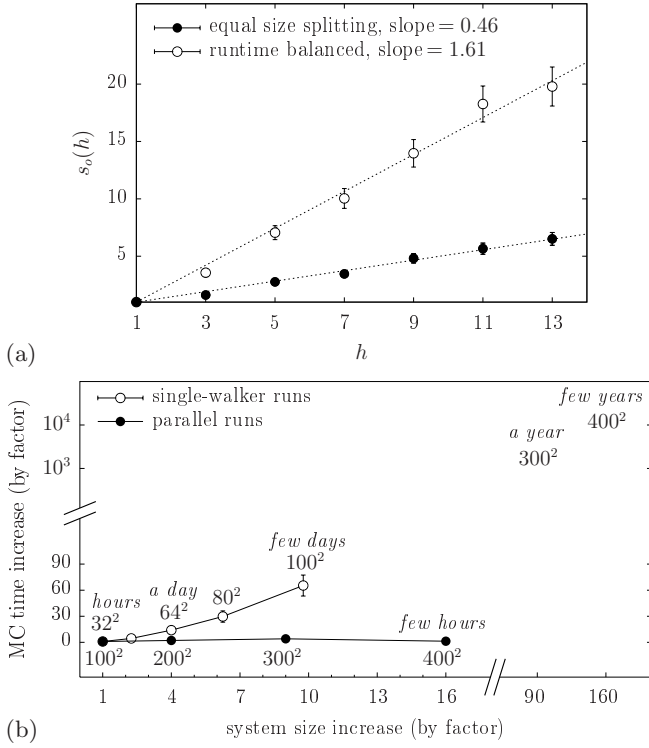


FIG. 12. Scaling results: (a) Speedups $s_o(h)$ for different numbers of energy windows h of equal size and overlap $o = 75\%$ (filled circles, cf. Fig. 2) and using run-time balanced energy splitting (open circles, cf. Fig. 1c), measured using the continuous lipid model. $h = 1$ corresponds to the single walker (one CPU) reference runs. The speedup is determined by the MC sweeps needed to complete the first WL iteration. Straight dotted lines are fits to the data. (b) Simulation-time increase for increasing system size for the 10-state Potts model. Results for single-walker runs (open symbols); parallel runs (filled symbols).

over the whole energy range. Second, there could be an algorithmic speedup through the introduction of the replica exchange move, an effect which is independent of the speedup from parallelization using multiple physical CPU cores. Indeed, timing experiments on a moderate-sized (48×48) Potts model using 15 energy windows revealed such an algorithmic speedup of up to roughly a factor of 2, depending on the frequency of the replica exchange move, compared to a scheme without replica exchanges (data not shown). Hence, not only does the replica exchange move between multiple energy windows reduce systematic errors in the WL simulations, it can potentially increase the algorithmic efficiency.

However, we stress that the actual performance and the relative importance of different mechanisms would vary from model to model; they are also related to the settings of the simulations. The speedup results presented above are merely examples and should not be taken as universal values for our framework.

In Fig. 12b we show the growth of simulation time with system size. We use the 2D Potts model here as

the scaling of the system is straightforward and the corresponding increase of the global energy range is known. We compare single-walker runs (open symbols) with parallel runs (filled symbols) where we increase the number of energy windows correspondingly with the increase in system size. Hence the size of individual energy windows and the overlap are fixed and we add windows as needed to cover the entire global energy range. By doing this we can keep the simulation time practically constant compared to the increase in simulation time for single-walker runs.

In summary, we find both strong and weak scaling: (i) for a given global energy range, a small number of additional energy windows can increase the speedup significantly, and (ii) system sizes can be extended without increasing the run time significantly by introducing additional individual energy windows.

V. SUMMARY AND PERSPECTIVES

We thoroughly investigated the properties of our recently introduced parallel framework for replica-exchange Wang-Landau sampling (REWL) [32]. The basic idea is to restrict individual WL walkers to small, but overlapping, energy windows and enable them to communicate with neighbors such that replica of the system can travel through the whole energy space. In contrast to traditional replica exchange Monte Carlo for which an unfortunate choice of temperatures leads to little overlap of probability distributions and effectively eliminates exchange of replicas, our scheme insures the possibility of replica exchange by fixing overlap of energy windows at the outset. We demonstrated the strength of the simulation framework via a sophisticated data analysis procedure to connect the resulting pieces of the density of states. We were able to reproduce very accurate and precise results of single-walker simulations, only much faster, and to facilitate simulations of systems which were not at all accessible before. By applying the proposed framework to qualitatively very different, challenging models, we showed that the method is generally applicable and robust. While it is possible to reduce the systematic errors by employing multiple walkers in an energy window, we also demonstrated that the proposed method shows both weak and strong scaling when increasing the number of computing cores. This parallel framework is much more efficient than single-walker WL sampling even on small hardware architectures like multi-core CPUs but, more importantly, it can be readily implemented on larger systems, potentially with $> 10^5$ processors, by making use of all lines of parallelization presented here.

The description of our method was held intentionally as simple and general as possible. Hence, despite its proven advantages in its current form, the procedure leaves much room for further optimization. Potential improvements are conceivable by fine-tuning energy win-

dows sizes and overlaps, frequency of replica-exchange moves, special treatment at window boundaries, etc. It is obvious, for example, that walkers in different equal-size energy windows will not proceed through the WL scheme at equal pace, i.e., the time it takes to fulfill the flatness criteria will greatly differ, particularly in the first iterations when the estimator for the density of states still differs significantly from the true DOS [32]. At that stage, the shape of the histogram $H(E)$ is not really flat. Hence, an optimal energy window size distribution would be based on the area under the respective local histograms such that the flatness criteria in each window will be fulfilled after approximately the same number of MC steps. On the other hand, in later stages of the iteration, when the walkers perform almost random walks through energy space, equal size energy windows might be indeed favorable. Eventually, one would like to end up

with self-tuning, variable-size energy windows such that all walkers proceed synchronously through the parallel WL scheme at all times.

ACKNOWLEDGMENTS

This work has been supported by the National Science Foundation under Grants DMR-0810223 and OCI-0904685. Y. W. Li was partly sponsored by the Office of Advanced Scientific Computing Research; U.S. Department of Energy. Part of the work was performed at the Oak Ridge Leadership Computing Facility at ORNL, which is managed by UT-Battelle, LLC under Contract No. De-AC05-00OR22725. Supercomputer time was provided by TACC under XSEDE grant PHY130009. LA-UR-13-29579 assigned.

-
- [1] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *J. Chem. Phys.* **21**, 1087 (1953).
 - [2] B. A. Berg and T. Neuhaus, *Phys. Lett. B* **267**, 249 (1991).
 - [3] B. A. Berg and T. Neuhaus, *Phys. Rev. Lett.* **68**, 9 (1992).
 - [4] A. P. Lyubartsev, A. A. Martsinovski, S. V. Shevkunov, and P. N. Vorontsov-Velyaminov, *J. Chem. Phys.* **96**, 1776 (1992).
 - [5] E. Marinari and G. Parisi, *Europhys. Lett.* **19**, 451 (1992).
 - [6] Y. Iba, *Int. J. Mod. Phys. C* **12**(5), 623 (2001).
 - [7] F. Wang and D. P. Landau, *Phys. Rev. Lett.* **86**, 2050 (2001).
 - [8] F. Wang and D. P. Landau, *Phys. Rev. E* **64**, 056101 (2001).
 - [9] C. Yamaguchi and Y. Okabe, *J. Phys. A: Math. Gen.* **34**, 8781 (2001).
 - [10] B. J. Schulz, K. Binder, and M. Müller, *Intl. J. Mod. Phys. C* **13**, 477 (2002).
 - [11] C. Yamaguchi and N. Kawashima, *Phys. Rev. E* **65**, 056710 (2002).
 - [12] N. Rathore and J. J. de Pablo, *J. Chem. Phys.* **116**, 7225 (2002).
 - [13] Q. Yan, R. Faller, and J. J. de Pablo, *J. Chem. Phys.* **116**, 8745 (2002).
 - [14] M. S. Shell, P. G. Debenedetti, and A. Z. Panagiotopoulos, *Phys. Rev. E* **66**, 056703 (2002).
 - [15] F. Calvo, *Mol. Phys.* **100**, 3421 (2002).
 - [16] M. Troyer, S. Wessel, and F. Alet, *Phys. Rev. Lett.* **90**, 120201 (2003).
 - [17] V. Mustonen and R. Rajesh, *J. Phys. A: Math. Gen.* **36**, 6651 (2003).
 - [18] F. Rampf, W. Paul, and K. Binder, *Europhys. Lett.* **70**, 628 (2005).
 - [19] C. Zhou, T. C. Schulthess, S. Torbrügge, and D. P. Landau, *Phys. Rev. Lett.* **96**, 120201 (2006).
 - [20] J. Luettmmer-Strathmann, F. Rampf, W. Paul, and K. Binder, *J. Chem. Phys.* **128**, 064903 (2008).
 - [21] M. P. Taylor, W. Paul, and K. Binder, *J. Chem. Phys.* **131**, 114907 (2009).
 - [22] T. Wüst and D. P. Landau, *Phys. Rev. Lett.* **102**, 178101 (2009).
 - [23] C. J. Geyer, in *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface*, ed. by E. M. Keramidas (Interface Foundation, Fairfax Station, VA, 1991), p. 156.
 - [24] K. Hukushima and K. Nemoto, *J. Phys. Soc. Jpn.* **65**, 1604 (1996).
 - [25] D. J. Earl and M. W. Deem, *Phys. Chem. Chem. Phys.* **7**, 3910 (2005).
 - [26] H. G. Katzgraber, S. Trebst, D. A. Huse, and M. Troyer, *J. Stat. Mech. (JSTAT)* **2006**, P03018 (2006).
 - [27] E. Bittner, A. Nußbaumer, and W. Janke, *Phys. Rev. Lett.* **101**, 130603 (2008).
 - [28] N. Rathore, T. A. Knotts, and J. J. de Pablo, *J. Chem. Phys.* **118**, 4285 (2003).
 - [29] M. O. Khan, G. Kennedy, and D. Y. C. Chan, *J. Comput. Chem.* **26**, 72 (2005).
 - [30] L. Zhan, *Comput. Phys. Commun.* **179**, 339 (2008).
 - [31] J. Yin and D. P. Landau, *Comput. Phys. Commun.* **183**, 1568 (2012).
 - [32] T. Vogel, Y. W. Li, T. Wüst and D. P. Landau, *Phys. Rev. Lett.* **110**, 210603 (2013).
 - [33] D. P. Landau, S.-H. Tsai, and M. Exler, *Am. J. Phys.* **72**, 1294 (2004).
 - [34] T. Wüst and D. P. Landau (to appear)
 - [35] T. Nogawa, N. Ito, and H. Watanabe, *Phys. Rev. E* **84**, 061107 (2011).
 - [36] C. Zhou and R. N. Bhatt, *Phys. Rev. E*, **72**, 025701(R) (2005).
 - [37] R. E. Belardinelli and V. D. Pereyra, *Phys. Rev. E* **75**, 046701 (2007).
 - [38] T. Vogel, Y. W. Li, T. Wüst, and D. P. Landau, *J. Phys.: Conf. Ser.* **487**, 012001 (2014).
 - [39] Y. W. Li, T. Vogel, T. Wüst, and D. P. Landau, *J. Phys.: Conf. Ser.* **510**, 012012 (2014).
 - [40] M. E. J. Newman und G. T. Barkema, *Monte Carlo methods in statistical physics*, Oxford University Press (Oxford, New York, 1999), and references therein.
 - [41] R. B. Potts, *Math. Proc. Camb. Phil. Soc.* **48**, 106 (1952).
 - [42] F. Y. Wu, *Rev. Mod. Phys.* **54**, 235 (1982).

- [43] W. Janke, Phys. Rev. B **47**, 14757 (1993).
- [44] R. J. Baxter, J. Phys. C: Solid State Phys. **6**, L445 (1973).
- [45] K. A. Dill, Biochemistry **24**, 1501 (1985).
- [46] M. Bachmann and W. Janke, Phys. Rev. E **73**, 020901(R) (2006).
- [47] R. Unger and J. Moulton, J. Mol. Biol. **231**, 75 (1993).
- [48] N. Lesh, M. Mitzenmacher, and S. Whitesides, RE-COMB '03 Proc. 7th Ann. Int. Conf. Res. Comp. Mol. Biol., pp. 188–195.
- [49] J. M. Deutsch, J. Chem. Phys. **106**, 8849 (1997).
- [50] T. Wüst and D. P. Landau, Comput. Phys. Commun. **179**, 124 (2008).
- [51] T. Wüst and D. P. Landau, J. Chem. Phys. **137**, 064903 (2012).
- [52] Y. W. Li, T. Wüst and D. P. Landau, Comput. Phys. Commun. **182**, 1896 (2011).
- [53] T. Wüst, Y. W. Li, and D. P. Landau, J. Stat. Phys. **144**, 638 (2011).
- [54] Y. W. Li, T. Wüst and D. P. Landau, Phys. Rev. E **87**, 012706 (2013).
- [55] R. Goetz and R. Lipowsky, J. Chem. Phys. **108**, 7397 (1998).
- [56] S. Fujiwara, T. Itoh, M. Hashimoto, and R. Horiuchi, J. Chem. Phys. **130**, 144901 (2009).
- [57] A. Seelig and J. Seelig, Biochem. **13**, 4839 (1974).
- [58] J. A. Aronovitz and M. J. Stephen, J. Phys. A: Math. Gen. **20**, 2539 (1987).
- [59] V. Blavatska and W. Janke, J. Chem. Phys. **133**, 184903 (2010).
- [60] T. Vogel, L. Gai, C. McCabe, P. T. Cummings, D. P. Landau, J. Chem. Phys. **139**, 054505 (2013).
- [61] H. K. Lee, Y. Okabe, and D. P. Landau, Comput. Phys. Commun. **175**, 36 (2006).
- [62] J. Zierenberg, M. Marenz, and W. Janke, Comput. Phys. Commun. **184**, 1155 (2013).
- [63] W. Bland, A. Bouteiller, T. Herault, J. Hursey, G. Bosilca, and J. J. Dongarra, in J. L. Träff, S. Benkner, and J. J. Dongarra (eds.) 19th EuroMPI, Proceedings. LNCS vol. 7490 (Springer, Berlin, 2012), pp. 193–203.